
Multi-relational Learning Using Weighted Tensor Decomposition with Modular Loss

Ben London, Theodoros Rekatsinas, Bert Huang, and Lise Getoor

University of Maryland

College Park, MD 20742, USA

{blondon, thodrek, bert, getoor}@cs.umd.edu

Abstract

We propose a modular framework for multi-relational learning via tensor decomposition. In our learning setting, the training data contains multiple types of relationships among a set of objects, which we represent by a sparse three-mode tensor. The goal is to predict the values of the missing entries. To do so, we model each relationship as a function of a linear combination of latent factors. We learn this latent representation by computing a low-rank tensor decomposition, using quasi-Newton optimization of a weighted objective function. Sparsity in the observed data is captured by the weighted objective, leading to improved accuracy when training data is limited. Exploiting sparsity also improves efficiency, potentially up to an order of magnitude over unweighted approaches. In addition, our framework accommodates arbitrary combinations of smooth, task-specific loss functions, making it better suited for learning different types of relations. For the typical cases of real-valued functions and binary relations, we propose several loss functions and derive the associated parameter gradients. We evaluate our method on synthetic and real data, showing significant improvements in both accuracy and scalability over related factorization techniques.

multi-relational data. We consider a scenario in which we are given a fixed set of objects, a set of relations and a small training set, sampled from the full set of all potential pairwise relationships; our goal is to predict the unobserved relationships. The relations we consider may be binary-, discrete ordinal- or real-valued functions of the object pairs; for the binary-valued relationships, the training labels include both positive and negative examples.

There has been a growing interest in tensor methods within machine learning, partially due to their natural representation of multi-relational data (Kashima et al., 2009). Many contributions (Dunlavy et al., 2006, 2011; Gao et al., 2011; Xiong et al., 2010) use the *canonical polyadic* (CP) decomposition, a generalization of singular value decomposition to tensors. Others (Bader et al., 2007) have proposed models based on *decomposition into directional components* (DEDICOM) (Harshman, 1978). We propose a similar decomposition (based on (Nickel et al., 2011)) which is more appropriate for multi-relational data, for reasons discussed in Section 3.2. Unlike these previous methods, we do not attempt to decompose the input tensor directly; rather, we explicitly model a mapping from the low-rank representation to the observed tensor, which is often better suited for prediction. For example, a binary relationship can be modeled as the sign of a latent representation; this gives the latent representation more freedom to increase the prediction margin, rather than reproduce $\{\pm 1\}$ exactly. In this respect, approaches like *maximum-margin matrix factorization* (MMMF) (Srebro et al., 2005b; Rennie and Srebro, 2005a) and DEDICOM can be viewed as specializations of our framework.

Our proposed method, *multi-relational weighted tensor decomposition* (see Section 3), assumes that the latent representation is determined by a linear combination of latent factors associated with each object. Learning these latent factors and their interactions in each relation thus becomes analogous to a weighted

1 Introduction

In network or relational data, one often finds multiple types of relations on a set of objects. For instance, in social networks, relationships between individuals may be personal, familial, or professional. We refer to this type of data as *multi-relational*. In this paper, we propose a tensor decomposition model for transduction on

tensor decomposition (described in Section 3.1 and illustrated in Figure 1). We formulate this decomposition as a nonlinear optimization problem (Section 3.3), which may incorporate any combination of smooth, task-specific loss functions. These task-specific loss functions allow simultaneous learning of various relation types, such as binary- and continuous-valued. By weighting the objective function, we are able to learn from limited observed (training) relationships without fitting the unobserved (testing) ones, improving both accuracy and efficiency. We demonstrate the effectiveness of our approach in Section 4, using both real and synthetic data experiments. Our results indicate that our approach is both more accurate and efficient than competing factorizations when training data is sparse.

2 Preliminaries

This section introduces our notation and defines the problem of multi-relational transduction.

We denote tensors and matrices using bold, uppercase letters; similarly, we use bold, lowercase letters to denote vectors. For a tensor \mathbf{X} , let $x_{i,j,k}$ denote the $(i, j)^{\text{th}}$ element of the k^{th} frontal slice. Denote by \mathbf{X}_k the matrix comprising the k^{th} frontal slice. We use \odot to denote the *Hadamard* (i.e., element-wise) product, $\text{tr}(\cdot)$ for the trace operator and $\|\cdot\|_F$ for the Frobenius norm. For a matrix \mathbf{V} and function f , let $\nabla_{\mathbf{V}} f$ denote the gradient of f with respect to \mathbf{V} .

Fix a set of m objects and a set of n relations.¹ To simplify our analysis, we assume that all relations are symmetric, though one can obtain an analogous derivation for asymmetric relations with only slightly more work. We are given a partially observed tensor $\mathbf{Y} \in \mathbb{R}^{m \times m \times n}$, in which each observed entry $y_{i,j,k}$ is a (possibly noisy) measurement of a relationship and each unobserved entry is set to a null value.² We are additionally given a nonnegative weighting tensor $\mathbf{W} \in \mathbb{R}^{+m \times m \times n}$, where each entry $w_{i,j,k} \in [0, 1]$ corresponds to a user-defined confidence, or certainty, in the value of $y_{i,j,k}$; if $y_{i,j,k}$ is unobserved, then $w_{i,j,k}$ is necessarily zero. The goal of multi-relational transduction in this tensor formulation is to infer the unobserved entries in \mathbf{Y} .

3 Proposed Method

This section introduces our proposed method, which we refer to as *multi-relational weighted tensor decom-*

position (MrWTD). We begin by describing our low-rank tensor representation of multi-relational data. We then define an optimization objective used to compute this representation and discuss how we solve the optimization.

3.1 Representation as Tensor Decomposition

Our fundamental assumption is that each relationship is equal to a mapping Φ_k applied to an element $x_{i,j,k}$ in an underlying low-rank tensor $\mathbf{X} \in \mathbb{R}^{m \times m \times n}$. Each Φ_k depends on the nature of the relation, and may differ across relations. For example, for binary relations in $\{\pm 1\}$, Φ_k is the *sign* function. We further assume that each \mathbf{X}_k can be factored as a rank- r decomposition

$$\mathbf{X}_k = \mathbf{A} \mathbf{R}_k \mathbf{A}^\top + b_k, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times r}$, $\mathbf{R}_k \in \mathbb{R}^{r \times r}$ and $b_k \in \mathbb{R}$. (Figure 1 illustrates this decomposition.) Note that there is a single \mathbf{A} matrix, but n instances of \mathbf{R}_k and b_k . Also note that we place no constraints on \mathbf{A} or \mathbf{R}_k ; the columns of \mathbf{A} need not be linearly independent, and \mathbf{R}_k need not be positive-semidefinite. To infer the values of the missing (or uncertain) entries, we predict each $y_{i,j,k}$ by computing $x_{i,j,k} = \mathbf{a}_i \mathbf{R}_k \mathbf{a}_j^\top + b_k$, where \mathbf{a}_i and \mathbf{a}_j are the i^{th} and j^{th} row vectors of \mathbf{A} , and then apply the appropriate mapping $\Phi_k(x_{i,j,k})$.

The entries of \mathbf{A} can be interpreted as the *global latent factors* of the objects, where the i^{th} row \mathbf{a}_i corresponds to the latent factors of object i . Each \mathbf{R}_k determines the interactions of \mathbf{A} in the k^{th} relation. Thus, each predicted relationship comes from a linear combination of the objects' latent factors. Because the latent factors are global, information propagates between relations during the decomposition, thus enabling collective learning. The addition of b_k accounts for distributional bias within each relation.

3.2 Related Models

Our tensor model is comparable to Harshman's DEDICOM (1978). Bader et al. (2007) applied the DEDICOM model to the task of temporal link prediction (in a single network), using the third mode as the time dimension. Recently, Nickel et al. (2011) proposed a relaxed DEDICOM, referred to as RESCAL, to solve several canonical multi-relational learning tasks. Of the previous approaches, our underlying decomposition is most similar to RESCAL, and Equation 1 would be identical to the RESCAL decomposition if not for the bias term. Beyond the decomposition, the key distinction is that RESCAL directly decomposes the input tensor, rather than modeling the mapping from \mathbf{X} to \mathbf{Y} . RESCAL also ignores the potential sparsity and uncertainty in the observations, whereas we explicitly

¹Here, we use the term *relation* loosely to include not only strict relations, for which relationships are either present or not, but also real-valued functions.

²For example, for binary-valued relations in $\{\pm 1\}$, the null value is 0.




Figure 1: For $k = 1, \dots, n$, each slice \mathbf{Y}_k of the input tensor is approximated by a function Φ_k of a low-rank decomposition $\mathbf{A}\mathbf{R}_k\mathbf{A}^\top + b_k$. The latent factors \mathbf{A} are common to all slices. Each \mathbf{R}_k determines the interactions of \mathbf{A} in the k^{th} relation, while b_k accounts for distributional bias.

model this. We demonstrate in Section 4 that our formulation produces more accurate predictions even when observed (training) data is limited.

Other tensor factorization models have been proposed for multi-relational data, though they typically use the CP decomposition (Dunlavy et al., 2006, 2011; Gao et al., 2011; Xiong et al., 2010). In the CP decomposition, each entry is the inner product of three vectors; this would be similar to our decomposition if each \mathbf{R}_k slice were constrained to be diagonal. The richer interactions of the relaxed DEDICOM and the global latent representation of the objects often make it better suited for multi-relational learning, as was corroborated empirically by Nickel et al. (2011).

3.3 Objective

To compute the decomposition in Equation 1, we minimize the following regularized objective:

$$f(\mathbf{A}, \mathbf{R}, \mathbf{b}) \triangleq \frac{\lambda}{2} \|\mathbf{A}\|_{\text{F}}^2 + \sum_{k=1}^n \frac{\lambda}{2} \|\mathbf{R}_k\|_{\text{F}}^2 + \text{tr}(\mathbf{W}_k(\ell_k(\mathbf{Y}_k, \mathbf{X}_k))^\top), \quad (2)$$

where $\lambda \geq 0$ is a regularization parameter, \mathbf{X}_k is computed by Equation 1, and ℓ_k is a loss function that is applied element-wise to the k^{th} slice. (For brevity, we use f to denote $f(\mathbf{A}, \mathbf{R}, \mathbf{b})$.) This ability to combine multiple loss functions is central to our approach, as the appropriate penalty depends on the mapping for each \mathbf{X}_k to \mathbf{Y}_k . Though most matrix and tensor decompositions focus on minimizing the quadratic loss (defined below), this criterion may not be optimal for certain prediction tasks (such as binary prediction). By explicitly making the loss function for each slice task-specific, our framework offers more flexibility than related techniques. The only requirement (due to our optimization method) is that the loss function is smooth.

It is important to note our use of L2 regularization. Regularization effectively controls the complexity of

the model and thereby reduces the possibility of over-fitting. This follows the traditional wisdom that “simpler” models will generalize better to unseen data—in this case, the unobserved tensor entries. The rank of the decomposition can also be seen as a complexity parameter, since higher ranks will better fit the observed data. However, after a certain point, increasing the rank has a diminishing effect, since the regularizer seeks to minimize the Frobenius norm of the decomposition. We explore the effect of the rank parameter empirically in Section 4.4.

To minimize Equation 2, we require the gradients of f w.r.t. \mathbf{A} , \mathbf{R}_k and b_k . Leveraging the symmetry of \mathbf{R}_k , we derive³ these as

$$\nabla_{\mathbf{A}} f = \lambda \mathbf{A} + \sum_{k=1}^n 2(\mathbf{W}_k \odot \nabla_{\mathbf{X}_k} \ell_k(\mathbf{Y}_k, \mathbf{X}_k)) \mathbf{A} \mathbf{R}_k^\top, \quad (3)$$

$$\nabla_{\mathbf{R}_k} f = \lambda \mathbf{R}_k + \mathbf{A}^\top (\mathbf{W}_k \odot \nabla_{\mathbf{X}_k} \ell_k(\mathbf{Y}_k, \mathbf{X}_k)) \mathbf{A}, \quad (4)$$

$$\nabla_{b_k} f = \text{tr}(\mathbf{W}_k (\nabla_{\mathbf{X}_k} \ell_k(\mathbf{Y}_k, \mathbf{X}_k))^\top), \quad (5)$$

where \odot denotes the *Hadamard* (i.e., element-wise) product, and $\nabla_{\mathbf{X}_k} \ell_k(\mathbf{Y}_k, \mathbf{X}_k)$ is the gradient of ℓ_k w.r.t. \mathbf{X}_k . Though this accommodates any differentiable loss function, we now present three that are applicable to many relational problems, and derive their corresponding loss gradients.

Quadratic Loss: The most common loss function used in matrix and tensor factorization is the *quadratic loss*, which we denote by $\ell^q(y, x) \triangleq \frac{1}{2}(y - x)^2$. Minimizing the quadratic loss corresponds to the setting in which each relationship is directly approximated by a linear combination of latent factors; i.e., Φ_k is the identity and $\mathbf{Y}_k \approx \mathbf{X}_k$. For this loss function, the loss gradient is simply $\nabla_{\mathbf{X}_k} \ell_k^q(\mathbf{Y}_k, \mathbf{X}_k) = (\mathbf{X}_k - \mathbf{Y}_k)$.

Smooth Hinge Loss: While the quadratic loss may be appropriate for learning real-valued functions, it

³Due to space restrictions, we state the gradients without their derivation.

is sometimes ill-suited for learning binary relations, which are essentially binary classifications. For binary classification, the goal is to complete a partially observed slice $\mathbf{Y}_k \in \{\pm 1\}^{m \times m}$. Recall that the mapping Φ_k is the sign function, and so $y_{i,j,k} \approx \text{sgn}(x_{i,j,k})$. Approximating $\{\pm 1\}$ with a quadratic penalty may yield a “small-margin” solution, since high-confidence predictions will push low-confidence predictions closer to the decision boundary. To get a “large-margin” solution, we use the *smooth hinge loss* (Rennie and Srebro, 2005a), $\ell^h(y, x) \triangleq h(yx)$, where

$$h(z) \triangleq \begin{cases} 1/2 - z & \text{if } z \leq 0, \\ (1 - z)^2/2 & \text{if } 0 < z < 1, \\ 0 & \text{if } z \geq 1. \end{cases}$$

Unlike the standard hinge loss, the smooth hinge is differentiable everywhere. To obtain closed-form gradients, we define tensors $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{m \times m \times n}$, where

$$p_{i,j,k} \triangleq \begin{cases} 1 & \text{if } 0 < y_{i,j,k}x_{i,j,k} < 1, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$q_{i,j,k} \triangleq \begin{cases} 1 & \text{if } y_{i,j,k}x_{i,j,k} < 1, \\ 0 & \text{otherwise.} \end{cases}$$

We can therefore express the smooth hinge as

$$h(y_{i,j,k}x_{i,j,k}) = (p_{i,j,k}x_{i,j,k}^2 - 2q_{i,j,k}y_{i,j,k}x_{i,j,k} + q_{i,j,k})/2,$$

which we can differentiate w.r.t. \mathbf{X}_k to obtain $\nabla_{\mathbf{X}_k} \ell_k^h(\mathbf{Y}_k, \mathbf{X}_k) = (\mathbf{P}_k \odot \mathbf{X}_k - \mathbf{Q}_k \odot \mathbf{Y}_k)$.

Logistic Loss: For binary relations, we can also use the *logistic loss* (Rennie and Srebro, 2005b), defined as $\ell^l(y, x) \triangleq \log(1 + e^{-yx})$. From a statistical perspective, this corresponds to the negative conditional log-likelihood of a logistic model. Note that this loss function also maximizes the binary prediction margin yx . The gradient of ℓ^l is easily derived as $\nabla_{\mathbf{X}_k} \ell_k^l(\mathbf{Y}_k, \mathbf{X}_k) = -\mathbf{Y}_k \odot \mathbf{Z}_k$, where $z_{i,j,k} \triangleq (1 + e^{yx})^{-1}$.

3.4 Weighting and Efficiency

The weighting tensor \mathbf{W} is a particularly important component of our framework. Without \mathbf{W} , the objective function would place equal importance on fitting both observed and unobserved values. If the observed tensor is very sparse (as it often is in real training data), this will result in fitting a large number of “phantom zeros”. The weighting tensor prevents this from happening by emphasizing only the observed (or certain) entries. We can thus train on a small number of observations without fitting the unobserved

data. This approach is similar to Acar et al.’s (2010), though their analysis is limited to the minimizing the quadratic loss for a CP decomposition.

Weighting the objective by \mathbf{W} also leads to an improvement in efficiency. When \mathbf{W} is sparse, the objective and gradient calculations are fairly lightweight, because any expression involving \mathbf{W} can be computed using sparse arithmetic. For instance, \mathbf{X}_k only appears in a Hadamard product with \mathbf{W}_k , so Equation 1 can be implemented as a sparse outer product, where we only compute $x_{i,j,k}$ for any nonzero $w_{i,j,k}$. In Equations 2–5, the only expressions that do not involve \mathbf{W} are the regularization terms. Thus, when \mathbf{W} has only c nonzero elements, the computational costs of these equations are $O(ncr + nmr^2)$. In contrast, methods that ignore the sparsity of the observed tensor take $O(nm^2r + nmr^2)$ time. Assuming that m^2 is the dominant term and that c grows much slower than m^2 (e.g., in natural networks, c is often $O(m)$), the sparse computation can be an order of magnitude faster.

Additionally, since \mathbf{W} can be real-valued (not just $\{0, 1\}$), we can adjust the entries to reduce the mistake penalty of certain examples. For instance, suppose an incorrect negative prediction is deemed more critical than an incorrect positive (as is often the case in medical diagnoses and certain link prediction tasks). One could multiplicatively increase the values of all $\{w_{i,j,k} : y_{i,j,k} = 1\}$ or, alternatively, decrease the values of all $\{w_{i,j,k} : y_{i,j,k} = -1\}$. This would effectively penalize false negatives more severely than false positives, encouraging the optimization to satisfy positive examples.

3.5 Optimization

To minimize the objective in Equation 2, we use *limited-memory Broyden-Fletcher-Goldfarb-Shanno* (L-BFGS) optimization. Since quasi-Newton methods, such as L-BFGS, avoid computing the Hessian, they are efficient for optimization problems involving many variables. All this requires is the objective function and the gradients in Equations 3–5. Since our optimization problem is non-convex, we are not guaranteed that L-BFGS will find the global minimum; in practice, however, the algorithm typically finds useful, though possibly local, minima.

To mitigate the possibility of finding local minima, we initialize the parameters using the eigendecomposition of each input slice, which is close to the desired factorization. This technique is similar to the initialization used by Bader et al. (2007) and Nickel et al. (2011), which have similar decompositions. For $k = 1, \dots, n$, let $\mathbf{\Lambda}_k \triangleq (\lambda_{1,k}, \dots, \lambda_{r,k})$ denote the r largest eigenvalues of \mathbf{Y}_k , and let $\mathbf{V}_k \triangleq (\mathbf{v}_{1,k}, \dots, \mathbf{v}_{r,k})$ denote their

corresponding eigenvectors. We initialize \mathbf{R}_k as a diagonal matrix with \mathbf{A}_k along the diagonal, and \mathbf{A} as the average of $\mathbf{V}_1, \dots, \mathbf{V}_n$. In practice, we find that this initialization converges faster, and often to a better solution, than random initialization.

Note that when the objective function uses only quadratic loss, one can compute the parameter updates using the *alternating simultaneous approximation, least squares and Newton* (ASALSAN) algorithm (Bader et al., 2007), which produces an approximate solution and has been shown to converge quickly. Since our objective may contain a heterogeneous mixture of loss functions—not all necessarily quadratic—we do not use ASALSAN. Moreover, we cannot use traditional convex programming techniques like *semidefinite programming* (SDP) because our objective is non-convex.

4 Experiments

In this section, we compare variants of MrWTD with RESCAL (Nickel et al., 2011), MMMF (Rennie and Srebro, 2005a) and *Bayesian probabilistic tensor factorization* (BPTF) (Xiong et al., 2010) in several experiments, using both real and synthetic data. The real data sources are kinship data from the Australian Alyawarra tribe, and two social interaction datasets from the MIT Media Lab. The comparisons highlight the critical advantages of MrWTD: namely, the ability to learn from limited training data, handle a mixture of learning objectives, transfer information across relations for collective learning, and exploit sparsity for improved efficiency.

To test the effect of the rank parameter, we run an experiment varying only the rank of the decomposition over a range of values. The results support our hypothesis that L2 regularization reduces the impact of the rank, effectively controlling the model complexity.

Finally, we perform a synthetic experiment to compare the running time of MrWTD to that of the above competing methods, demonstrating the significant scalability gains provided by exploiting sparsity.

To conserve space, certain figures and tables are provided in the supplementary material (Appendix A).

4.1 Compared Methods

To evaluate the performance of various loss functions, we compare several variants of MrWTD. The variant named MrWTD-Q uses the quadratic loss for all relations, regardless of their type. MrWTD-H and MrWTD-L use the quadratic loss for real-valued slices and the smooth hinge or logistic loss, respectively, for binary slices.

The RESCAL model approximates each slice of the input tensor as $\mathbf{Y}_k \approx \mathbf{A}\mathbf{R}_k\mathbf{A}^\top$. In (Nickel et al., 2011), binary relationships are represented by $\{0, 1\}$. Unfortunately, since RESCAL does not account for missing data, unobserved relationships are simply treated as negative examples. In order to distinguish between (un)observed relationships and negative examples, we use $\{\pm 1\}$ for observed data and zeros elsewhere. In our experiments, we find that this modification improves RESCAL’s performance over the original method. Since RESCAL uses the quadratic loss uniformly, it uses ASALSAN to compute the decomposition, with L2 regularization on \mathbf{A} and \mathbf{R}_k .

MMMF is a tool for matrix reconstruction and, as such, is not designed for multi-relational data. That said, we can use it to reconstruct each slice of the tensor individually. Like MrWTD, MMMF approximates a binary input \mathbf{Y} using the sign of a rank- r matrix decomposition, $\mathbf{Y} \approx \text{sgn}(\mathbf{U}\mathbf{V}^\top)$, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{m \times r}$. The “fast” variant of the algorithm (Rennie and Srebro, 2005a) adds a bias term and uses the smooth hinge loss. The optimization objective is very similar to ours, but with different gradients, due to the decomposition. Our implementation of fast MMMF differs from that of Rennie and Srebro (2005a) only in the way we solve the optimization (using L-BFGS, rather than conjugate gradient descent) and the fact that the input is assumed to be symmetric.

Because the synthetic data generator (described in Section 4.2) matches our decomposition and is slightly different than that of traditional MMMF, it is somewhat unfair to compare traditional MMMF to MrWTD. We therefore run a variant of MrWTD that decomposes each slice separately instead of jointly, using a separate \mathbf{A}_k . This is meant to equalize the discrepancy in the decomposition, while isolating the deficiencies of non-collective learning. We refer to this model as MMMF+.

BPTF is a fully Bayesian interpretation of the CP tensor factorization, originally designed for temporal prediction.⁴ We compare it to MrWTD to investigate the benefits and drawbacks of the Bayesian approach. BPTF assumes that all latent factors are sampled from a Gaussian distribution. The only user-defined properties are the rank of the decomposition and the hyperparameters. The parameters and latent factors are estimated using Gibbs sampling. One benefit of the Bayesian approach is that it avoids the model selection problem, which in our case is the choice of regulariza-

⁴Sutskever et al. (2009) propose another fully Bayesian algorithm, *Bayesian tensor factorization* (BTF), whose decomposition is very similar to ours, though their framework only supports the quadratic loss. We were unable to compare MrWTD to this method at the time of submission.

tion parameters.⁵ This can have a pronounced effect when training data is limited, which makes proper regularization critical. However, Gibbs sampling is computationally expensive, since it requires many iterations of sampling to converge to an accurate estimate. We analyze this trade-off between accuracy and efficiency in Section 4.5. Additionally, BPTF only supports the quadratic loss, since it has a natural probabilistic interpretation as the Gaussian likelihood and makes the model conjugate, making Gibbs sampling easier. No such interpretation exists for the (smooth) hinge loss, and the logistic loss has no conjugate prior.

We implement all of the above methods in MATLAB, using a third-party implementation of L-BFGS⁶, and the authors’ implementation of BPTF⁷.

4.2 Synthetic Data Experiments

To generate the synthetic data, we start by computing a low-rank tensor $\hat{\mathbf{X}} \in \mathbb{R}^{m \times m \times n}$ as $\hat{\mathbf{X}}_k \leftarrow \hat{\mathbf{A}} \hat{\mathbf{R}}_k \hat{\mathbf{A}}^\top + \mathbf{E}_k$, for $k = 1, \dots, n$, where $\hat{\mathbf{A}} \in \mathbb{R}^{m \times r}$ and $\hat{\mathbf{R}}_k \in \mathbb{R}^{r \times r}$ are sampled from a normal distribution, and $\mathbf{E}_k \in \mathbb{R}^{m \times m}$ is low-level, normally-distributed noise. For the first experiment, we construct $n = 3$ binary relations (i.e., slices), over $m = 500$ objects, using rank $r = 10$. We refer to this dataset as Binary Synthetic. To generate a binary tensor $\mathbf{Y} \in \{\pm 1\}^{m \times m \times n}$, we round the values of $\hat{\mathbf{X}}$ using the 90th percentile of its values as a threshold. This produces a heavy skew towards the negative class, as is typical in real multi-relational data. For the second experiment, we construct one binary relation and one real-valued relation, again over 500 objects, with rank 10. We normalize the real-valued relation such that the standard deviation is 1.0, giving it roughly the same scale as the binary slices. We refer to this dataset as Mixed Synthetic.

We evaluate over training sizes $t \in [3, 25]$ percent, averaging the results over 20 runs per size. In each run, we sample a random $t \cdot \binom{m}{2}$ pairs (and their symmetric counterparts) from each slice to use as the training set, and let the remaining pairs comprise the test set. We then hold out a random 25% from the training set as a validation set for a regularization parameter search, where we search over the range $[10^{-3}, 10^3]$ in logarithmic increments. For Binary Synthetic, we select the optimal parameter λ^* that maximizes the area under the precision-recall curve (AUPRC), averaged over all slices; for Mixed Synthetic, we maximize the harmonic mean of the AUPRC of the first slice and one minus the mean-squared error (MSE) of the second. We then

retrain on the full training set using λ^* and evaluate on the test set. For BPTF, we run Gibbs sampling for 200 iterations.

The results of the synthetic data experiments are given in Figure 2, reported as average AUPRC and MSE over 20 runs. On Binary Synthetic, MrWTD-L achieves a statistically significant⁸ lift over the competing methods for training sizes 5% and up, and all three variants showing significant lift for 10% and above. We attribute these results to two primary advantages: the weighted objective function, with its mixture of task-specific loss functions, and the global latent factors. As discussed in Section 3.5, the weighted objective is necessary for exploiting small amounts of observed (i.e., training) data, without fitting the unobserved entries. Since RESCAL treats all entries as observed, it tends to fit the unobserved entries in sparsely populated tensors. Furthermore, though MMMF and MMMF+ use the same large-margin technique as MrWTD-H and MrWTD-L, they do not perform collective learning, since the latent factors are specific to each slice. In MrWTD, information from one slice is propagated to the others via the global latent factors. Note that BPTF and MrWTD-Q perform significantly worse than the large-margin loss variants of MrWTD for sizes 10% and above, illustrating that the quadratic loss is not always appropriate for binary data. On Mixed Synthetic, MrWTD’s improvement over RESCAL and MMMF, for both slices, is statistically significantly for all training sizes. MMMF+ is competitive with MrWTD on the real-valued slice, with significant lift for training sizes 3, 5%; yet its performance deteriorates on the binary slice, for all training sizes, since it is not able to transfer information between slices. BPTF is also competitive with MrWTD on the real-valued slice, and the binary slice for smaller training sizes, but falls slightly behind on the higher sizes.

4.3 Real Data Experiments

We evaluate on several real multi-relational datasets. The first dataset consists of kinship data from the Australian Alyawarra tribe, as recorded by Denham and White (2005). This data has previously been used by Kemp et al. (2006) for multi-relational link prediction. The data contains $m = 104$ tribe members and $n = 23$ types of kinship (binary) relations.⁹ In total, the dataset includes 125,580 related pairs. This yields a tensor $\mathbf{Y} \in \{\pm 1\}^{104 \times 104 \times 23}$.

The remaining datasets come from MIT’s Human Dy-

⁵As the authors claim, the effect of tuning the hyperparameter priors is minimal.

⁶www.di.ens.fr/~mschmidt/Software/minFunc.html

⁷www.cs.cmu.edu/~lxiong/bptf/bptf.html

⁸We measure statistical significance in all experiments using a 2-sample t-test with rejection threshold 0.05.

⁹The original data contains 26 relations, but relations 24–26 are extremely sparse, exhibiting fewer than 6 instances, so we omit them.




Figure 2: Results of the synthetic data experiments (discussed in Section 4.2). The horizontal axis indicates the size of the training data (in percentage of the tensor); in (a) and (b), the vertical axis indicates the area under the precision-recall curve; in (c), the vertical axis shows the mean squared error (MSE). All scores are averaged over 20 runs. The top to bottom arrangement of the legend corresponds to a left to right arrangement in each group. See Table 2 in the appendix for the full results, including standard deviations.

namics Laboratory. Both consist of human interaction data from students, faculty, and staff working on the MIT campus, recorded by a mobile phone application. From the first dataset, named Reality Mining (Eagle et al., 2009), we use the survey-annotated network, consisting of $n = 3$ types of binary relationships annotated by the subjects: friendship, in-lab interaction and out-of-lab interaction. These relationships are measured between $m = 94$ participants, providing a total of 13,395 related pairs. The resulting tensor is $\mathbf{Y} \in \{\pm 1\}^{94 \times 94 \times 3}$. From the second dataset, named Social Evolution (Dong et al., 2011), we use the survey-annotated network, as well as several interaction relations derived from sensor data, resulting in $n = 8$ binary relations with 16,101 related pairs. The five surveyed relations are: close friendship, biweekly social interaction, political discussion, two types of social media interaction. The three derived relations are computed from: voice calls, SMS messaging and proximity. We binarize this data by a simple indicator of whether the given type of interaction occurred. In this case, the number of users is $m = 84$, results in a tensor $\mathbf{Y} \in \{\pm 1\}^{84 \times 84 \times 8}$.

For these experiments, we use the same methodology as the synthetic experiments, with rank $r = 20$. The results are also given in Figure 3. The three variants of MrWTD and BPTF achieve significant lift over RESCAL, MMMF and MMMF+ in nearly all experiments. MrWTD has a statistically significant advantage over the other methods for most training ratios on the Kinship data, while BPTF has an advantage on the Social Evolution data. Yet, as we show in the following section, MrWTD’s estimation takes a small fraction of BPTF’s running time. We therefore achieve results that are comparable to Bayesian methods in far

less time. We refer the reader to Table 1 in the appendix for the complete set of results.

4.4 Rank Experiment

To measure the effect of the rank parameter on the performance of each algorithm, we rerun the Social Evolution experiment, varying $r = \{5, 10, 20, 40\}$ and keeping the training ratio is fixed at 25%. The results of this experiment are displayed in Figure 5, in the appendix. There is a small increase in AUC from $r = 5$ to $r = 10$, which is expected, since 5 is relatively low. However, we find that the effect of the rank is minimal for $r \geq 10$; the standard deviation across all runs in this range is < 0.02 for each algorithm. This supports our hypothesis that, beyond a certain threshold, the regularizer is the primary controller of model complexity.

4.5 Timing Experiment

Finally, we measure the running time of each of the above tensor methods to better understand their scalability in scenarios where training data is limited. We create a sequence of synthetic datasets (using the technique in Section 4.2), each with $n = 3$ binary slices, for sizes $m = \{500, 1000, 2000, 4000, 8000\}$. For training, we use a random 10% of the tensor. We compare the smooth hinge loss variant of MrWTD, RESCAL and BPTF, using predefined regularization and hyperparameters. We run these experiments on a machine with two 6-core Intel® Xeon® X5650 processors, running at 2.66 GHz, and 48 GB of RAM.

The timing results, averaged over 10 runs per problem size, are shown in Figure 4. BPTF takes considerably more time than the others, due to its Gibbs sampling




Figure 3: Results of the real data experiments (discussed in Section 4.3). The horizontal axis indicates the size of the training data (in percentage of the tensor); the vertical axis indicates the area under the precision-recall curve, averaged over 20 runs. The top to bottom arrangement of the legend corresponds to a left to right arrangement in each group. See Table 3 in the appendix for the full results, including standard deviations.

estimation. Note that we could not run BPTF on the two largest problem sizes, due to out-of-memory exceptions. This illustrates the tradeoff between accuracy and efficiency in using Bayesian methods; one can reduce running time by reducing the number of iterations, but this would also affect the accuracy of the estimation. Due to the efficient, closed-form updates of the ASALSAN algorithm, RESCAL is the fastest for small problem sizes. However, MrWTD is significantly faster as the problem size grows. This is because RESCAL’s objective function treats all tensor entries with equal importance, whereas MrWTD’s weighted objective only requires the predictions of the observed entries, thus allowing us to skip prediction on the test data during estimation.

5 Conclusion

In this paper, we present a modular framework for multi-relational learning via tensor decomposition. The decomposition we use provides an intuitive interpretation for the multi-relational domain, where objects have global latent representations and relationships are determined by a function of their linear combinations. We show that the global latent representations enable information to transfer between relation types during model estimation. Further, we demonstrate that our framework’s weighted objective and support for multiple loss functions improves accuracy over similar models. Finally, we show that our method exploits the sparsity of limited training data to achieve an order of magnitude speedup over unweighted methods.

We plan to extend MrWTD to be able to learn from large-scale data by adapting hashing methods from matrix factorization literature (Karatzoglou et al.,




Figure 4: Results of the timing experiment (discussed in Section 4.5). The horizontal axis indicates the size of a synthetic dataset, measured by the number of objects m ; the vertical axis indicates the running time (in seconds), averaged over 10 runs. For each dataset, we use 10% for training. Note that BPTF could not run on sizes $\{4000, 8000\}$ due to runtime exceptions.

2010). We would also like to compare our method to Sutskever et al.’s BTF algorithm (2009), to further investigate the benefit of the Bayesian approach. We also intend to analyze the theoretical properties of our framework, such as generalization error, using existing learning theory literature (Srebro et al., 2005a; Cortes et al., 2008; El-Yaniv and Pechyony, 2009).

Acknowledgements

This work was partially supported by NSF CAREER grant 0746930 and NSF grant IIS1218488.

References

- E. Acar, D. Dunlavy, T. Kolda, and M. Mørup. Scalable tensor factorizations with missing data. In *Proc. of the 2010 SIAM International Conf. on Data Mining (SDM)*, 2010.
- B. Bader, R. Harshman, and T. Kolda. Temporal analysis of semantic graphs using ASALSAN. In *Proc. of the 7th IEEE International Conf. on Data Mining (ICDM)*, 2007.
- C. Cortes, M. Mohri, D. Pechyony, and A. Rastogi. Stability of transductive regression algorithms. In *Proc. of the 25th International Conf. on Machine Learning (ICML)*, 2008.
- W. Denham and D. White. Multiple measures of Alyawarra kinship. *Field Methods*, 17(1), 2005.
- W. Dong, B. Lepri, and A. Pentland. Modeling the co-evolution of behaviors and social relationships using mobile phone data. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia*, MUM '11, pages 134–143, New York, NY, USA, 2011. ACM.
- D. Dunlavy, T. Kolda, and W. Kegelmeyer. Multilinear algebra for analyzing data with multiple linkages. Technical Report, 2006.
- D. Dunlavy, T. Kolda, and E. Acar. Temporal link prediction using matrix and tensor factorizations. *ACM Trans. on Knowledge Discovery from Data*, 5(2), 2011.
- N. Eagle, A. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. *Proc. of the National Academy of Sciences*, 106(36), 2009.
- R. El-Yaniv and D. Pechyony. Transductive Rademacher complexity and its applications. *J. Artificial Intelligence Research (JAIR)*, 35, 2009.
- S. Gao, L. Denoyer, and P. Gallinari. Link pattern prediction with tensor decomposition in multi-relational networks. In *IEEE Symposium on Comp. Intell. and Data Mining*, 2011.
- R. Harshman. Models for analysis of asymmetrical relationships. In *First Joint Meeting of the Psychometric Society and the Society for Mathematical Psychology*, 1978.
- A. Karatzoglou, A. Smola, and M. Weimer. Collaborative filtering on a budget. In *Proc. of the 13th International Conf. on Artificial Intelligence and Statistics*, 2010.
- H. Kashima, T. Kato, Y. Yamanishi, M. Sugiyama, and K. Tsuda. Link propagation: a fast semi-supervised learning algorithm for link prediction. In *SIAM International Conference on Data Mining (SDM)*, 2009.
- C. Kemp, J. Tenenbaum, T. Griffiths, T. Yamada, and N. Ueda. Learning systems of concepts with an infinite relational model. In *Proc. of the 21st National Conf. on Artificial Intelligence*, 2006.
- M. Nickel, V. Tresp, and H. Kriegel. A three-way model for collective learning on multi-relational data. In *Proc. of the 28th International Conf. on Machine Learning (ICML)*, 2011.
- J. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *In Proc. of the 22nd International Conf. on Machine Learning (ICML)*, 2005a.
- J. Rennie and N. Srebro. Loss functions for preference levels: regression with discrete ordered labels. In *IJCAI Multidisciplinary Workshop on Adv. in Preference Handling*, 2005b.
- N. Srebro, N. Alon, and T. Jaakkola. Generalization error bounds for collaborative prediction with low-rank matrices. In *Advances in Neural Information Processing Systems 17*. 2005a.
- N. Srebro, J. Rennie, and T. Jaakkola. Maximum-margin matrix factorization. In *Advances in Neural Information Processing Systems 17*. 2005b.
- I. Sutskever, R. Salakhutdinov, and J. Tenenbaum. Modelling relational data using Bayesian clustered tensor factorization. In *Advances in Neural Information Processing Systems 22*, pages 1821–1828. 2009.
- L. Xiong, X. Chen, T. Huang, J. Schneider, and J. Carbonell. Temporal collaborative filtering with Bayesian probabilistic tensor factorization. In *SIAM International Conference on Data Mining (SDM)*, 2010.

A Supplementary Material

Here we report additional results from the experiments discussed in Section 4. Figure 5 shows the results of the rank experiment (Section 4.4). Table 1 shows the results of the timing experiment (Section 4.5). In Table 2 and Table 3, we list the full results of the synthetic (Section 4.2) and real data (Section 4.3) experiments.




Figure 5: Results of the rank data experiment (discussed in Section 4.4), using the Social Evolution dataset. The horizontal axis indicates the rank of the decomposition; the vertical axis indicates the area under the precision-recall curve, averaged over 20 runs. We use a random 25% of the tensor for training data on each run. There is a slight increase from $r = 5$ to $r = 10$, but less than 0.02 standard deviation (per algorithm) for $r \geq 10$, supporting our hypothesis that regularization controls model complexity.

Table 1: Results of the timing experiment (discussed in Section 4.5). The first column indicates the size of a synthetic dataset, measured by the number of objects m ; the remaining columns indicate the running time (in seconds), averaged over 10 runs, with the associated standard deviations. For each dataset, we use 10% for training. Note that BPTF could not run on sizes $\{4000, 8000\}$ due to runtime exceptions.

m	RESCAL		BPTF		MrWTD-H	
500	0.21	(0.05)	1.52	(0.05)	0.59	(0.01)
1000	0.67	(0.10)	6.51	(0.42)	1.35	(0.37)
2000	3.39	(0.35)	25.47	(1.57)	5.25	(0.05)
4000	32.96	(5.02)	–	(–)	20.33	(0.80)
8000	332.77	(40.88)	–	(–)	80.18	(2.44)

Table 2: Results of the synthetic data experiments (discussed in Section 4.2). The first column indicates the dataset and score type. (For Mixed Synthetic, we provide two row groups to display the scores of the binary- and real-valued slices.) The second column indicates the amount of training data (in percentage of the tensor). We evaluate on three variants of MrWTD: -Q uses the quadratic loss for all slices; -H and -L use the quadratic loss for real-valued slices, but use the smooth hinge and logistic losses, respectively, for binary slices. For Binary Synthetic and the first slice of Mixed Synthetic, we report area under the precision-recall curve (AUPRC); for the second slice of Mixed Synthetic, we report mean squared error (MSE). Standard deviations are listed in parentheses. All scores are averaged over 20 runs. Bold scores are statistically tied with the best score in each row.

Data	Tr %	RESCAL	MMMF	MMMF+	BPTF	MrWTD-Q	MrWTD-H	MrWTD-L
Binary Synth (AUPRC)	3	0.14 (.00)	0.13 (.04)	0.13 (.00)	0.15 (.01)	0.13 (.01)	0.12 (.01)	0.13 (.01)
	5	0.16 (.01)	0.16 (.03)	0.17 (.01)	0.27 (.03)	0.23 (.03)	0.20 (.08)	0.32 (.03)
	10	0.21 (.03)	0.24 (.04)	0.36 (.01)	0.50 (.01)	0.55 (.02)	0.70 (.02)	0.71 (.02)
	15	0.38 (.03)	0.33 (.01)	0.51 (.01)	0.59 (.01)	0.66 (.02)	0.83 (.00)	0.84 (.00)
	20	0.49 (.02)	0.37 (.01)	0.61 (.01)	0.63 (.01)	0.70 (.02)	0.89 (.00)	0.90 (.00)
	25	0.57 (.02)	0.40 (.01)	0.69 (.01)	0.66 (.01)	0.73 (.02)	0.92 (.00)	0.93 (.00)
Binary Synth (AUPRC)	3	0.14 (.01)	0.11 (.03)	0.13 (.01)	0.26 (.02)	0.27 (.03)	0.27 (.03)	0.26 (.04)
	5	0.18 (.01)	0.13 (.01)	0.18 (.03)	0.37 (.02)	0.36 (.03)	0.37 (.03)	0.36 (.04)
	10	0.27 (.02)	0.30 (.02)	0.35 (.03)	0.57 (.02)	0.55 (.03)	0.55 (.03)	0.54 (.04)
	15	0.33 (.01)	0.25 (.08)	0.48 (.01)	0.65 (.01)	0.67 (.04)	0.67 (.02)	0.67 (.02)
	20	0.39 (.02)	0.20 (.03)	0.57 (.02)	0.70 (.01)	0.74 (.03)	0.74 (.02)	0.75 (.02)
	25	0.43 (.01)	0.24 (.06)	0.66 (.01)	0.71 (.01)	0.74 (.02)	0.74 (.02)	0.76 (.02)
Mixed Synth (MSE)	3	0.99 (.01)	0.99 (.01)	0.43 (.08)	0.48 (.03)	0.44 (.04)	0.44 (.04)	0.45 (.03)
	5	0.97 (.01)	0.96 (.02)	0.12 (.07)	0.24 (.02)	0.19 (.02)	0.19 (.01)	0.19 (.01)
	10	0.90 (.00)	0.92 (.01)	0.03 (.01)	0.06 (.01)	0.05 (.04)	0.04 (.00)	0.05 (.04)
	15	0.83 (.00)	0.90 (.02)	0.02 (.00)	0.03 (.00)	0.03 (.02)	0.03 (.02)	0.04 (.02)
	20	0.77 (.01)	0.84 (.00)	0.02 (.01)	0.02 (.00)	0.02 (.01)	0.03 (.01)	0.03 (.01)
	25	0.70 (.01)	0.81 (.00)	0.01 (.00)	0.01 (.00)	0.01 (.00)	0.01 (.00)	0.01 (.00)

Table 3: Results of the real data experiments (discussed in Section 4.3). Standard deviations are listed in parentheses. All scores are averaged over 20 runs. Bold scores are statistically tied with the best score in each row.

Data	Tr %	RESCAL	MMMF	MMMF+	BPTF	MrWTD-Q	MrWTD-H	MrWTD-L
Kinship (AUPRC)	3	0.08 (.00)	0.08 (.00)	0.08 (.00)	0.10 (.03)	0.10 (.04)	0.09 (.01)	0.11 (.03)
	5	0.08 (.01)	0.09 (.00)	0.09 (.00)	0.26 (.03)	0.33 (.02)	0.31 (.04)	0.33 (.02)
	10	0.13 (.01)	0.11 (.00)	0.15 (.01)	0.49 (.02)	0.44 (.13)	0.46 (.03)	0.48 (.02)
	15	0.18 (.01)	0.14 (.01)	0.21 (.01)	0.57 (.01)	0.61 (.01)	0.42 (.24)	0.57 (.12)
	20	0.28 (.01)	0.16 (.01)	0.27 (.01)	0.59 (.01)	0.65 (.02)	0.65 (.02)	0.63 (.01)
	25	0.34 (.01)	0.18 (.00)	0.33 (.02)	0.61 (.01)	0.68 (.01)	0.70 (.01)	0.69 (.01)
Reality (AUPRC)	3	0.09 (.01)	0.09 (.01)	0.09 (.01)	0.11 (.01)	0.09 (.01)	0.09 (.01)	0.09 (.01)
	5	0.09 (.01)	0.12 (.08)	0.10 (.01)	0.13 (.02)	0.11 (.02)	0.11 (.02)	0.11 (.02)
	10	0.10 (.01)	0.13 (.03)	0.13 (.02)	0.20 (.03)	0.17 (.03)	0.17 (.03)	0.18 (.04)
	15	0.11 (.01)	0.19 (.03)	0.17 (.03)	0.24 (.03)	0.27 (.06)	0.23 (.06)	0.23 (.08)
	20	0.13 (.01)	0.22 (.06)	0.21 (.04)	0.29 (.03)	0.32 (.04)	0.27 (.06)	0.29 (.08)
	25	0.14 (.01)	0.26 (.03)	0.25 (.05)	0.31 (.04)	0.34 (.02)	0.31 (.05)	0.34 (.04)
Social (AUPRC)	3	0.34 (.03)	0.29 (.02)	0.30 (.01)	0.41 (.01)	0.32 (.01)	0.32 (.01)	0.32 (.01)
	5	0.35 (.03)	0.30 (.01)	0.31 (.01)	0.45 (.02)	0.35 (.03)	0.36 (.04)	0.36 (.03)
	10	0.36 (.01)	0.33 (.01)	0.36 (.01)	0.51 (.01)	0.50 (.02)	0.47 (.03)	0.43 (.06)
	15	0.45 (.01)	0.38 (.04)	0.41 (.02)	0.56 (.01)	0.54 (.02)	0.54 (.02)	0.54 (.03)
	20	0.52 (.02)	0.41 (.02)	0.44 (.02)	0.60 (.02)	0.57 (.01)	0.57 (.02)	0.58 (.02)
	25	0.56 (.01)	0.44 (.01)	0.47 (.01)	0.63 (.01)	0.60 (.01)	0.60 (.01)	0.62 (.01)